

MULTICAST PEERING

Background of the Invention

An Internet is a packet switched network of computers and local networks consisting of nodes, which can be computers or networks of computers, routers,
5 and data transmission lines, with the routers being used to route packets over data transmission lines towards the intended recipient. All information on an Internet is conveyed encapsulated in packets, which consist of a header (containing routing and other information) and a body for containing data.

Generally data transfers on an Internet consist of more information than can be
10 conveyed in one packet, with a data transmission thereby consisting of a set of related packets being sent from one source to one receiver. A data transmission consisting of time ordered data, such as with an audio or video broadcast, is called a stream. Data transmissions on an Internet uses Internet Protocol (IP) standards based on protocols standardized by the Internet Engineering Task
15 Force (IETF). Each node and router reachable in an Internet is assigned an Internet address, and routers maintain information about how to reach any of a variety of Internet addresses. A transmission of packets from one router to the next in the chain lending to the final destination is called a "hop".

Unless a source and a receiver are directly connected, Internet
20 transmissions of data will pass through at least one, and in general many, routers between leaving the source and arriving at the receiver. These routers can use one of several routing protocols, which describe both the creation and storage in any particular router of information about how packets should be forwarded to

reach any particular destination, a means of sharing this information between routers, and other details about the forwarding of packets. Routing protocols also generally include adjustable parameters, describing such features as possible data transmission rates and the timing of various actions, which can be set by the owner or operator of the router based on various considerations. The choice of a particular set of parameter values in a given routing protocol is referred to as the router configuration. In addition, routers from specific vendors may or may not support all routing protocols, or all possible choices of router configuration.

IP data transmissions can be one to one (or unicasting), one to all (or broadcasting), one to many or many to many, with the last two possibilities being described as multicasting. Multicasting is more specifically a means of the sending of packets from one or more sources to one or more receivers, in such as way as only one copy of each packet is required to leave any source and the packets are multiplied as required by routers in the Internet to reach the desired receivers. Unicasting, multicasting and broadcasting of packets are all performed using specific routing protocols, and one router may use different routing protocols for each of these different transmission methods.

Broadcasting of packets is used for signaling and notification only, being explicitly confined to only the neighboring nodes and routers of the transmitters, with routers being forbidden to forward broadcast packets to other routers. Actual data transfers on an Internet thereby are conducted using unicasting or multicasting.

Multicast data transfers on an Internet are transmitted from a particular source, with a related set of transmissions being called a multicast group. There must be at least one source per group, and any particular source can only belong to one particular multicast group.

5 Any given node on an Internet can be a source for one or more groups; the number of sources that a node can support being restricted only by hardware or software limitations at the node, although each source in a group must have a unique Internet address.

10 In order to perform multicasting on an Internet, it is necessary to construct a multicast tree. In general, each source for each group must have a separate tree, but it is possible for sources for one group to share part or all of their trees, creating a shared tree. There is no central control in multicast routing; a tree is constructed by the routers and consists of a table maintained by each router, with said table containing the name of the group, a list of sources (for source specific
15 trees), the direction of the source (i.e., the location of the source, or of the next router along the tree to the source), and the direction to any receivers of the group traffic, with each receiver being assumed to be interested in all of the source transmissions for some group.

A1
20 The existence of these tables within a router is called router state. Unlike the case in unicasting, multicast transmission or reception changes the router state, through the additions of new multicast groups, sources, or receivers. Changes in router state can be a major cause of resource expenditure by a

network, with excessively frequent changes, or excessively large router tables, the potential to seriously degrade network performance.

A2 Multicast routing protocols are classified as "sparse mode" or "dense mode". In sparse mode, reception of a multicast transmission by a receiver is accomplished by a multicast join, which is a message sent from the receiver to the nearest router (the so called "first hop router"), requesting the transmission. If the router is already part of the multicast tree and is already receiving the transmission, then the transmission is simply routed to the new receiver. If not, the router sends a join message to the next router in the chain going to either the source (if known) or a rendezvous point (RP, also called a Core), if the location of the source is not known. The join request travels towards the source or the RP, either a router is reached that is already receiving the multicast transmissions, or until the source or the RP is reached. In sparse mode a receiver stops receiving a multicast transmission (i.e., leaves the multicast group), by sending a "prune" message to the first hop router, which then ceases forwarding transmissions to the receiver. Multicast state in the routers is always subject to timers, and required periodic refreshing to remain valid; because of this it is also possible to stop receiving multicast transmissions by remaining silent, and thus to "time out". In either case, each router in the tree, if it is no longer forwarding the multicast transmission to any receiver, will itself send a prune message to next router in the tree to be removed from the tree entirely.

In dense mode multicasting, which is an older and less capable technology, multicast packets are initially flooded to all possible receivers (the

A3> "flood" stage), which then have to explicitly prune themselves if they are not interested in receiving the transmissions (the "prune" stage). This "flood and prune" technique, although technically easy to implement, and is used on small Internets, or small subsets of large Internets, is not suited for deployment on
5 arbitrarily large Internets due to the geometrical multiplication of the data transmissions required during the initial flood stage.

A4> Although all of the routers in an Internet can use the same routing protocols with identical router configurations, in general routers in an Internet are owned and / or controlled by a number of entities, and it is common for there to
10 be different routing protocols and routing configurations chosen for business or technical reasons by different router operators in an Internet. As it can be difficult or impossible to transmit packets between routers using different protocols, or with the same protocol but with different configurations, it is common to divide an Internet into so-called Autonomous Systems, with each Autonomous System
15 being a set of routers, networks and nodes managed as one administrative unit, using either one routing protocol and configuration, or a compatible set of routing protocols and / or configurations, i.e., one routing policy. In general, the routers in a given Autonomous System are given sufficient information to be able to route a packet to any Internet address contained within the Autonomous System.

20 Unicast transmissions of data between Autonomous Systems are done through the use of specially chosen routers known as Border Routers, or BRs, with a Border Gateway Protocol (BGP) to facilitate the exchange of unicast routing information between different Autonomous Systems. Using these

protocols, a BR in one Autonomous System can discover whether a node address exists in another Autonomous System and details on unicast routing to that other node address in the other Autonomous System.

In general, routing of packets in an Internet is subject to business
5 arrangements and must be accounted and paid for. Since the boundaries
between Autonomous Systems generally coincide with a boundary between
different Service Providers, unicast transmissions over Border Routers are
generally accounted for and used as part of the calculations of the payments
owed by one service provider to another. Since the operators of two Autonomous
10 Systems can both account for the traffic at any Border Routers between the two
systems, each can audit the unicast transmissions between their Autonomous
Systems, and there is no need, in the unicast case, for intrusive audit
arrangements between competitive Service Providers. The set of arrangements
allowing for unicast transmissions between independent Service Providers is
15 called unicast peering, or simply peering.

The large scale use of multicasting on Internets with more than one
Autonomous Systems is currently hindered by the technical and business
difficulties of conveying multicast transmissions over the boundaries between
Autonomous Systems.

20 There is currently a lack of multicast peering on commercial Internets,
because of business and commercial problems associated with such multicast
peering. In multicasting, a single packet stream might cross a Border Router from
an Autonomous System One (AS1) into an Autonomous System Two (AS2), and

AS > there be multiplied into many separate streams. Accounting for this multicast traffic, so that AS1 can properly pay AS2 for the work entailed by this transmission, would thus require detailed knowledge of the internal traffic within AS2, and this might reveal proprietary information about AS2. An audit of this
5 accounting could not be done without intrusive monitoring of conditions within AS2, which would also reveal sensitive and proprietary information about the workings of AS1, AS2 .

There are various technical problems at present with inter domain multicast routing. Multicast transmissions are sent to a multicast group, which
10 consist of one or more receivers, from a source. If multicast traffic is allowed freely from the border routers of AS1 to those of AS2, there is no mechanism to prevent any source in AS1 from sending traffic to the multicast group in AS2. Unauthorized sources in AS1 could thus flood multicast group members in AS2 with unwanted packets, possibly disrupting the reception of intended data
15 transmissions, thereby constituting a Denial of Service Attack on AS2.

A technical solution under development to solve some, but not all, of the problems associated with multicasting to multiple Autonomous Systems is called single source multicasting (SSM, or also PIM-Source only, or PIM-SO). In SSM, each multicast group is allowed to have only one, specific, source, associated
20 with a specific Internet address. This alleviates the problems associated with unauthorized sources flooding an Autonomous System with multicast packets, however, it does not solve the business problems associated with accounting and managing multicast traffic across multiple Autonomous Systems. SSM is

ALP also efficient for solutions where there are multiple multicast sources emanating from one IP address, in that there has to be a separate multicast tree maintained for each such source. It is also possible that there would be three or more Autonomous Systems involved in a multicast transmission, say AS1, AS2 and AS3. In this situation, all sources might be located in AS1 and all receivers in AS3, but AS2 might be essential in the construction of the multicast tree between AS1 and AS3. AS2 is thus forced with performing work for which it has no customers, and thus no commercial reason to perform. This problem is called a "third party dependency" in the literature.

Summary of the Invention

The solution to the lack of multicast peering in commercial Internets that is the subject of the present invention is to perform multicast peering with a trusted third party. In this solution, the trusted third party has a connection into both Autonomous System One and Autonomous System Two. Multicast streams pass from the trusted third party into both Autonomous Systems independently. There is thus no need for any sharing of information between Autonomous System One and Autonomous System Two, with any information sharing taking place only between each Autonomous System and the trusted third party.

The inventive trusted third party solution also solves various technical problems associated with multicast transmissions between different Autonomous Systems.

In the present invention, for example, the problem associated with maintaining separate multicast trees for each source, is solved because multicast trees can be shared by all such sources, reducing the amount of information that has to be maintained by each router in the multicast tree.

5 In the trusted third party solution, any multicast packets entering into two Autonomous Systems, AS1 and AS2, come only from the Trusted Third Party (TTP), which can control any multicast transmissions into each Autonomous System. There is no possibility of a malicious or unintended transfer of multicast transmissions from an arbitrary location in another Autonomous System. If for
10 some reason the amount of multicast transmissions from TTP becomes too large, the trusted third party transmissions come from a known location, and can be restricted or terminated by the network operators.

Third party dependencies are avoided in the trusted third party solution, as there are contractual relationships between TTP and each Autonomous System
15 that participates in the multicast transmissions.

In a particularly advantageous embodiment of the invention, Staggered Erasure Correction (SEC) which is designed to deal with the time-correlated nature of actual packet loss, is used by TTP.

A7 > In yet another particularly advantageous embodiment of the invention,
20 TTP is set up to prevent unauthorized sources from transmitting data on the TTPN through the use of a Multicast Firewall.

Brief Description of the Drawings

FIG. 1 is a block diagram which illustrates an implementation of a business model for multicast peering in accordance with an embodiment of the present invention.

5

Detailed Description of the Preferred Embodiments of the Invention

In a preferred implementation of the business model in accordance with the present invention, a dedicated-to-many transmission of multicast packets from one location to unlimited numbers of receivers is created using the Protocol Independent Multicast – Sparse Mode version 2 (PIM-SM v.2), with a co-located and restricted Rendezvous Point (RP), located at the TTP facility. Transmissions will consist of a number of separate channels, with each channel itself consisting of a number of separate sub-channels. In initial operations, each TTP channel will consist of 4 separate sub-channels, in order to provide the necessary information for the Staggered Erasure Correction (SEC), which is described in detail subsequently, and to provide for auxiliary text and computer control information. The delivery of information in a TTP channel shall be referred to as a stream, where it is implicitly understood that this stream consists of separate sub-streams, one for each of the sub-channels. The facilities and contractual relationships needed to perform these transmissions will, in aggregate, be called the Trusted Third Party Network (TTPN).

10

15

20

FIG. 1 is a block diagram which illustrates a trusted third party solution, wherein any multicast packets entering into two Autonomous Systems, AS1 and AS2, come from the Trusted Third Party (TTP) rather than as cross-border traffic between AS1 and AS2. In such an implementation, TTP can control any

5 multicast transmissions into each Autonomous System, AS1 and AS2.

Multicast groups will be created for each separate TTP stream, with each multicast group consisting of a number of separate sources, one source for each sub-stream created for each sub-channel. TTP will host a group specific RP for these groups, and only for these groups, and these groups will not be advertised

10 for other RP's in the Autonomous Systems with which TTP has a contractual relationship. In PIM-SM v.2, multicast packet delivery starts out using a shared tree rooted at the RP, which is a Shortest Path Tree (SPT) for the RP, but is not an SPT for a arbitrarily located source. In PIM-SM v.2, when data transmissions cross a specified threshold data rate for a particular group, that group is

15 transferred to a SPT for each source in that group. This transfer is done at the router where the SPT for the RP and the SPT for the source diverge. In the TTPN, the sources are at the RP, so this transfer will not occur. This offers the advantage of both the Shortest Path Tree and the reduction in router state provided by use of a Core Based Tree.

20 A8 > It is generally thought that not switching to a SPT for each source is not good practice, and, indeed, in the commodity Internet the general practice is to set the threshold to zero, so that the transfer to the source based SPT occurs immediately. This is done to avoid having a "hot spot" at the RP, which would

have to handle the routing for all sources in the group, and to keep the multicast traffic exclusively on SPTs. In the case of the one-to-many static transmissions on the TTPN, these benefits are illusory, as all traffic will use a SPT, and because some router would have to be the first hop router for the TTPN traffic,
5 and would thus have to accommodate all of the TTPN traffic.

In yet another embodiment of the invention, TTP is set up to prevent unauthorized sources from transmitting data on the TTPN through the use of a Multicast Firewall as follows.

- 10 The only way to initiate, terminate or modify a multicast transmission in PIM-SM v.2 is through
- reception of a PIM register message at the RP (used to register a new source and begin transmissions)
 - reception of a PIM register-stop message at the RP (used to stop transmissions from a source)
 - 15 . Reception of a IGMP v.2 membership report which mentions a new source or group
 - . Reception of PIM join / leave messages (these apply only to multicast receivers, not to sources).
 - . Reception of unexpected multicast traffic,

20

The multicast firewall will protect the TTPN RP router from unauthorized state changes by

- Rejection of PIM register and register-stop message from outside the TTPN facility,

- Rejection of IGMP v.2 membership reports which refer to multicast groups or sources not used by the TTPN.

5 A9 > . Reception of PIM join / leave messages which refer to multicast groups or sources not used by the TTPN.

- Multicast traffic from outside (i.e., all multicast traffic flows outward only).

The use of this multicast firewall will thus protect against unauthorized
10 transmissions from the TTPN facilities (including transmissions from elsewhere that would then immediately be switched to a SPT not using the TTPN RP), as well as unauthorized termination of the TTPN transmissions.

In yet another embodiment of the invention, TTP will employ Staggered Erasure Correction (SEC) in order to ensure quality transmissions. SEC is
15 described as follows.

The Internet is a "best-effort" packet transmission service, where there is no guarantee that a particular packet will be delivered, or that packet streams will be delivered in the order transmitted. There are various protocols, such as TCP, that attempt to guarantee packet delivery through the use of retransmissions of
20 lost packets upon the request of the receiver. These methods do not have an obvious analogue in multicasting. At present, attempts to provide for guaranteed packet delivery using multicasting have not been very successful. In accordance

with this embodiment of the invention, TTPN does not use any form of reliable multicasting but instead achieves acceptable performance using SEC. Crucial to the design of a particular implementation of SEC is the Maximum Dropout Period (MDP), the longest interval of packet dropouts that can be tolerated.

5 On the commodity Internet, at times of network congestion, packet loss can reach 10% or more on transcontinental routes. Unless provision is made for packet loss, each packet loss (or "dropped packet") represents a break in the audio or video signal. Unlike the case in the unicast Transmission Control Protocol (TCP), which allows for packet recovery, the multicast transmission
10 protocols do not attempt to retransmit dropped packets. Packet loss tends to be "bursty", with periods of zero or low loss interspersed with periods of high, or even total, packet loss. In mathematical language, packet loss tends to be time correlated : if a packet is lost at a specific time , the probability that the next packet will be lost is higher than the probability that an arbitrarily selected packet,
15 distant in time from that specific time, will be lost. This increased probability of additional packet losses after the loss of a packet declines with time from the event. Over a given duration, D , the packet loss can thus be statistically described in terms of a mean loss rate during that duration, ϵ , and a correlation time, τ , the interval over which the correlation of packet losses drops to an
20 insignificantly low value. (Note that D must be larger than both τ and the duration between packets for this statistical description to be appropriate.)

The time-correlated nature of packet loss will tend to defeat erasure protection methods based on the simultaneous retransmission of a given stream.

A simple erasure correction scheme, for example, would be to simply transmit each packet twice in a row, so that if a given packet was dropped, it could be simply replaced with the following packet. Suppose that the mean packet loss rate at a given time is 1 %, or $\epsilon = 0.01$. If packet losses were not time correlated,

5 then the probability that two subsequent packets would be lost is ϵ^2 , or one chance in ten thousand in this example, which might be acceptable for many applications. Since packet loss is time-correlated this scheme would be much less effective in practice. The chances that two subsequent packets would be lost might be very high, even for a mean loss rate of 1%. If the chances that, given
10 one packet is lost, the subsequent one is lost too is 50%, then the actual stream drop-out rate resulting from this scheme is 0.5 % in this example, only a modest improvement in performance at the cost of doubling the transmitted bandwidth.

The Staggered Erasure Correction (SEC) was designed to deal with the time-correlated nature of actual packet loss. Multiple copies of the same
15 information are sent staggered in time, with the stagger interval being selected to be larger than a typical value for the packet loss correlation time.

A10 > In that case, the separate streams will have uncorrelated packet losses, and, for a given mean loss rate, ϵ , and N separate staggered streams, the total loss rate will be ϵ^N . In actual practice, a mean loss rate 10 % (or $\epsilon = 0.1$)
20 represents a high rate of packet loss, and a typical value for packet loss correlation time would be 1 second. Given these parameters, the following table describes the expected performance of a SEC system :

SEC Performance : $\varepsilon = 0.1$

	Number of Staggered Channels	Probability of Packet Loss	Mean Time Between Drop-outs
5	1	10%	10 seconds
	2	1%	100 seconds
	3	0.1 %	1000 seconds

SEC Performance : $\varepsilon = 0.01$

	Number of Staggered Channels	Probability of Packet Loss	Mean Time Between Drop-outs
10			
15	1	1%	100 seconds
	2	0.01%	2.8 hours
	3	0.0001 %	12 days

All 20 with a typical drop-out duration being 1 second. A SEC with three staggered streams is sufficient to reduce drop outs to the level of a few per hour under fairly extreme conditions of packet loss, and to near zero at other times, which was the MTN design goal.

Staggered Erasure Correction : Sub-streams

A12

In the most straight-forward SEC implementation, each staggered stream would be a full copy of the original data. The above SEC implementation, with three staggered streams, would reduce drop-outs to a few per hour, but at the cost of tripling the bandwidth required. In many cases, however, such as for entertainment and most audio and video transmissions, a degraded copy of the data stream, at a reduced bandwidth, may be an acceptable replacement for the full data rate. In audio entertainment, for example, using psycho-acoustic compression, while a bandwidth of 160 kilobits per second is required to give full sound quality, a bandwidth of 64 kilobits per second still provides acceptable stereo sound reproduction at most times, and a bandwidth of 32 kilobits per second is marginally acceptable in monaural sound reproduction. Since a stereo sound reproduction can be sent as two monaural reproductions, the staggered channels used in SEC can be one full rate channel (the main stream), which would, in conditions of no packet loss, be the source of the sound reproduction, plus two monaural channels (the sub-streams). In the case a full rate channel packet was dropped, it would be replaced by the two monaural channels, used to reproduce the stereo audio stream, while it would take the loss of both the main rate stream and one of the two sub-streams before the reproduction quality dropped to monaural. This scheme provides the same protection against dropouts as the full channel reproduction SEC, but at a cost of only a 40% increase in bandwidth required, compared to the 200% increase required by the

full channel reproduction SEC.

The following table describes the expected durations of degraded signal quality in a period of high packet loss.

5 SEC Reproduction Quality : $\varepsilon = 0.1$

	Number of Staggered Channels	Probability of Packet Loss	Mean Time Between Drop-outs
10	1	10%	10 seconds
	2	1%	100 seconds
	3	0.1 %	1000 seconds
15	Sound Quality		Seconds per hour
	Full rate Stereo		3240
	Reduced rate Stereo		324
	Monaural		32
20	Drop-out		4

Staggered Erasure Correction : Recombination

The SEC scheme requires recombination to recreate the original stream. In the case of a stream ordered in time, such as an audio or video stream, each packet conveys (part or all) of the signal for a given interval of time. 5
Recombination is facilitated if these periods are the same or commensurate for each of the SEC streams. At the receiver, packets are decoded and time ordered in a separate queue for each stream. Suppose that each packet represents the signal for a time t_i . In the receiver, there will thus be a queue of packets for t_i , t_{i-1} , 10 t_{i-2} , etc, with the possibility of missing packets from any queue. When it is time to reproduce the signal for time t_i , the receiver examines each queue in turn, selecting either the first queue with a packet for that time, if the queue is for a full-rate channel, or packets for as many reduced rate channels as are available. Only in the case that every queue was missing a packet for that time is there a 15 dropout. The reconstituted stream can then be sent to other software, such as a video or audio player, for further processing; this subsequent software need not know the details of the SEC recombination process.

20 ~~ATTP~~ can implement SEC in its initial broadcasts through, for example, transmission of four separate sub-streams :

- ~~- The Main sub-channel (M-channel), a joint normal stereo MP3 encoding at 160 kilobits per second (kbps), transmitted MDP / 2~~

A13

seconds in advance of real time (i.e., the time at which the transmissions are intended to be played).

- The Immediate sub-channel (I-channel), a mono MP3 encoding at 32 kbps, transmitted 1 second in advance of real time.
- 5 - The Delayed sub-channel (D-channel), a mono MP3 encoding at 32 kbps, transmitted MDP seconds in advance of real time.
- The Text sub-channel (T-channel), transmitting the ASCII text required for the advertising crawl bar and any required control information, transmitted MDP / 2 seconds in advance of real time at 5 kbps.

10

In this example, in normal full rate operations, the M, I, D and T channels will each be separate sources of a multicast group, and all will be received by the full rate service. Since every separate multicast source is charged for, the number of multicast groups will have to be kept to a minimum. Until IGMP version 3 support

15 is available in the network, there will have to be separate multicast groups for

- G₁ : Full rate service: the M channel, totaling 160 kbps.

A14

~~G₂ : Lower rate service: the D channels, totaling 32 kbps.~~

- G₃ : Basic service: the I-channel and T-channel, at 39 kbps.

These multicast groups can be used to support, for example, the following

20 services

- S₁ : Full rate service: the G₁, G₂ and G₃ groups, totaling 229 kbps.

- S_2 : Lower rate service: the G_1 and G_3 groups only, totaling 197 kbps.
- S_3 : ISDN service: using the G_2 and G_3 groups only, totaling 69 kbps.
- S_4 : Dial-up service: using the G_3 group only, totaling 39 kbps.

5 In the primary usage, the receiver will select the appropriate service for its connection bandwidth.

When initially joining the transmission, the I channel information is used to provide low rate audio broadcast until all of the other channels have been received.

10 When changing channels, the new G_3 group is joined immediately, and playback starts after 3 seconds. At that time the other groups, if any, in the service are joined.

The SEC will require, for stereo broadcasts, encoding the following information, where L is the Left stereo channel feed and R is the Right stereo

15 channel feed :

- The M channel is the usual joint stereo
- The D channel is encoded as $L + R$ in mono.
- The I channel is encoded as $L - R$ in mono.

20

When stereo is made available from the D and I channels, then the following assignments must be made :

- Left = $(D + I) / 2$
- 5 - Right = $(D - I) / 2$

In case of a mono recording, the left and right channels will both be the mono feed.

10 The SEC is implemented through use of common time slices for each MP3 frame (which will consist of more than one Internet packet in general). The receiver will order the incoming frames from each group in a random access buffer. The following playback order shall be observed :

- 15 - 1.) For a time slice with M, D and I frames, the M frame shall be played.
- 2.) For a time slice with M and I frames, the M frame shall be played.
- 3.) For a time slice with M and D frames, the M frame shall be played.
- 4.) For a time slice with D and I frames only, normal stereo shall be played using the above decoding
- 20 - 5.) If only the D or I frame is available, mono shall be played based on the available frame.

In all cases the playback shall be in real time (i.e., 1 second after the encoding time of the I channel).

A16

- 5 These separate groups can be used to implement Receiver-Based Congestion avoidance (WBCA) based on the amount of time the primary channel has to be replaced by lower rate information. If this occurs for more than WBCA Threshold Ratio proportion of the time in a WBCA threshold interval, the receiver shall implement WBCA. The default values for the WBCA Threshold Ratio is 50%, and for WBCA Interval is 5 minutes.

A17

- 10 In PIM-SM v.2 a multicast group leave is supposed to be completed in no more than 3 seconds. The MTN receivers, if receiving Group S_i , with $i < 4$, can execute receiver based congestion avoidance by going from S_i to service S_{i+1} by:

- 1.) Leaving the group, G_i , with the lowest index i in S_i
- 15 - 2.) (For service S_2 only) Waiting 3 seconds, using the previously stored SEC information to make up for any packet loss
- 3.) (For service S_2 only) Joining group G_2 .

- 20 This process can be repeated as needed. After a period of time, the Congestion Avoidance Wait Period (CAWP, with a default of 5 minutes), then an attempt is made to restore the previous service by going from S_{i+1} to service S_i by ;

- 1.) (For service S_3 only) Leaving group G_2 .
- 2.) (For service S_3 only) Waiting 3 seconds, using the previously stored SEC information to make up for any packet loss
- 3.) Joining group G_j required for service S_i

5

A18 > While various implementations of the inventive system and model for multicast peering, including Staggered Erasure Protection and Multicast Firewall, have been described in detail, a skilled artisan will readily appreciate that numerous other implementations, particularly those where establishing a
10 multicast transmission is desired, are possible without departing from the spirit of the invention.